



# 802.1D/Q Compliance and Spatial Reuse


Bridging Adhoc Spatial Reuse Subteam

Li Mo

William Dai

John Coulter

Robert Castellano



5/2/2002, Version 2.3

bah\_spat\_01.pdf



## Outline



- Objective
  - Definitions
  - Overall Spatial Reuse strategy
  - Major issues Spatial Reuse needs to address
  - Spatial Reuse Requirements
- Spatial Reuse Control Sublayer
  - Example
  - Support for Basic Bridging Compliance Model
- Conclusions



## Presentation Objectives



- Define a set of procedures to address the 802.17 PAR requirement for spatial reuse while satisfying the compatibility requirements with 802.1D/Q
  - Preserve the 802.1 bridging relay ISS/EISS
  - Define a generalized form of spatial reuse that can be supported by all forms of stations on the ring.
  - Investigate the scenarios/issues and ensure the integrity of 802.1D/Q compliant bridges when devices on the ring invoke spatial reuse.
  - Define the set of transmit/receive rules stations must adhere to when stripping frames from the ring.
  - Define how spatial reuse would coexist with the basic bridging compliance model (if required).



## Definitions



- **Spatial Reuse** – Frames are stripped from the ring based on the frame's destination address. Not all the stations see a copy of the frame on the ring.
- **Flooded Frame** – Frame is stripped from the ring upon completion of being transmitted to every station on the ring. Every station on the ring sees the flooded frame.
- **Station ID** – Identifier used in the 802.17 frame that uniquely corresponds with a station on the ring. The source station ID, indicates the originating station on the ring. The destination station ID indicates the intended recipient of a frame.
- **MAC Address** – The address used to transmit a frame from one end station to another on either the ring or the bridged network.
- **MAC/FID** – MAC address / Forwarding Information Database. Defines the filtering method for a given VLAN ID. FID entries can be defined as IVL or SVL. In IVL each VID defines a separate filtering database. In SVL, multiple VID's share the same filtering database. It is important to have consistent filtering rules when MAC addresses are associated with more than 1 VLAN.



## Overall Strategy



- Stations shall employ a station mapping algorithm as part of ringlet selection which explicitly identifies the intended receiver which is to strip the packet.
  - The station mapping algorithm resides in the transmitting station which maps destination end station MAC addresses/FID to ring station ID addresses.
  - The transmitting station shall adhere to a set of ringlet and station selection policies to ensure that frames are not black-holed, misordered, or persistently flooded.
  - Transit stations which are not the intended receiver do not copy a unicast frame to their MAC client. Transit stations shall only copy broadcast frames. All unicast frames must be destination stripped by the intended receiver. (A special *flood\_only* signal in the frame is needed to indicate to bridges when they should/should not copy transit frames, if basic bridging compliance mode of operation continues to be supported.)



## Major issues to address with spatial reuse



- Packet Misordering
  - Packet Misordering may occur if a station's ringlet selection alternates transmission between the two ringlets for a given MAC/FID pair.
- Black Holing
  - Black Holing may occur if there is a topology change due to Spanning Tree or an end station moves from residing behind one bridge attached to the ring vs. another.

### Persistent Bridge Flooding

- Persistent Bridge Flooding may occur if a bridge relay only sees a unidirectional stream between two end stations, and never has the opportunity to learn the intended destination address. The bridge relay always floods on all its other network ports. Asymmetrical transmissions on the ring can cause the filtering mechanism in the bridge to break down.




## Packet Misordering / 1



Host X and Host Y are equal distance on CW / CCW paths. Host X may select CW path to transmit to Host Y. Host Y may select CW path to transmit to Host X

Host Y sees Host X transmission on CW path and begins transmitting on CCW path. Host X sees Host Y transmission on CW path and begins transmitting on CCW path.

Every time Host X and Host Y changes paths, results in packet misordering. There may be cases where Host X and Host Y continuously oscillate between CW and CCW ringlets leading to continuous packet misordering



**Solution** – Allow asymmetric transmission between Host X and Host Y. Or, define some kind of precedence which ensures only 1 of the two stations shall change ringlets.




## Packet Misordering / 2



Host X and Host Y communicate through two disparate VLANs – VLAN 1 and VLAN 2. If Host X does not distinguish between VLAN1 and VLAN2, then Host X first learns MAC Y on the CCW ringlet, then on the CW ringlet. As Host X changes transmission between CW and CCW ringlets, packet misordering occurs.

**Solution** – Host MAC X must be able to perform IVL MAC/FID ringlet selection or allow asymmetric transmission between MAC X and MAC Y.








## Black Holing / 1



Host X and Host Y communicate through two disparate VLANs – VLAN 1 (through Bridge A) and VLAN 2 (through Bridge B). If Host X does not distinguish between VLAN1 and VLAN2, then if Host X transmits all Host Y traffic via Bridge A, all VLAN 2 traffic is black-holed, and vice-versa.

**Solution** – Host X must be able to perform IVL MAC/FID ringlet selection in its station mapping table. All Host Y, VLAN 1 traffic must be forward via bridge A. All Host Y, VLAN 2 traffic must be forwarded via bridge B.






## Black Holing / 2



Host X and Host Y are communicate through two bridges (bridge A and B) which are connected the same ring. Spanning Tree protocol is used to break the loop formed by Host Y, Bridge A, and Bridge B by blocking the path between B and Y. Host MAC X forwards traffic to Host Y via Bridge A. If the link to Bridge A fails, STP establishes a new path between Host Y and the ring through Bridge B. If Host X continues to forward to Y via Bridge A, then all of X's frames sent to Y are black-holed.



**Solution** – Spanning Tree Protocol shall issue a TCN message which is intended to speed the aging of stale entries in the FDB. The RPR MAC shall support an RPR TCN control message that relays the TCN information in a standardized format to all stations on the ring. The bridge normally transmitting the TCN also transmits the RPR TCN. This removes dependency on vendor specific TCN formats.



## Persistent Flooding / 1




TBs flood all frames X transmits to Y

**Problem:** Bridges do not properly learn host addresses due to Destination Stripping by X and Y. Bridges continuously flood all unicast packets to all other LAN networks, defeating the function of the filtering bridge.

**Solution:** Intermediate bridges should not copy transit unicast traffic destined for another station to their bridging relay. Intermediate bridges shall only copy broadcast or unknown. Known unicast traffic is not flooded on the ring or throughout the adjoining bridged LAN networks. Only broadcast/unknown traffic is flooded.

TBs flood all frames Y transmits to X






## Persistent Flooding /2



TBs flood all frames X transmits to Y

**Problem:** Bridges do not properly learn host addresses due to Destination Stripping by X. Bridges continuously flood all unicast packets on the ring, defeating the function of the filtering bridge.

**Solution:** Intermediate bridges should not copy transit unicast traffic destined for another station to their bridging relay. Intermediate bridges shall only copy broadcast or unknown. Known unicast traffic is not flooded on the ring or throughout the adjoining bridged LAN networks. Only broadcast/unknown traffic is flooded.





## 802.1D/Q Spatial Reuse Requirements / 1



Stations implementing spatial reuse shall implement the

### Spatial Reuse Sublayer (SRCs)

- SRCs shall support the following ringlet selection and transmission rules.
- SRCs shall support RPR TCN control messages

Frame reception and stripping rules apply to all stations.

## Ringlet Selection Rules

- Stations shall maintain a station mapping database which maps MAC\_DA/FID to a destination station ID / ringlet ID. The FID policy (SVL / IVL) should be consistently configured for all stations on the ring.
- Stations shall explicitly encode a Destination Station ID for all known unicast. Broadcast/Unknown traffic shall be flooded to all stations on the ring using a special DSID identifier (0xFF for example).
- Station mapping database shall be updated based on learning MAC\_SA/FID w/ source station ID. Ringlet ID may be associated through learning, precedence algorithm or through topology database. Care must be taken in selecting the ringlet ID in order to avoid packet misordering. The ringlet selection rules do allow for known unicast traffic to have asymmetrical transmit/receive paths.
- An aging timer is required for station table mapping entries. Entries which expire are removed from the table.



## 802.1D/Q Spatial Reuse Requirements / 3



- Frame Reception / Stripping Rules
  - Stations shall receive all frames with a DSID that explicitly matches the station's ID.
  - Stations shall receive all multicast frames.
  - Stations shall receive all broadcast frames. (DSID set to 0xFF).
  - Unicast frames with a DSID that does not match the station ID are not passed to the station's bridging relay (transit only). Station's may choose to receive these frames to facilitate MAC/FID learning.
  - Frames having an explicit destination station ID shall only be received and stripped by the intended destination station. Frames having the explicit broadcast identifier shall be flooded, and stripped after all stations receive the frame.
  - All received frames MAC\_SA/FID/SSID are passed to the station mapping table for learning and to the station's bridge relay or MAC client for forwarding.

## Support of Spanning Tree Topology Changes

- The RPR mapping tables need to support accelerated aging in response to Spanning Tree Topology changes in the network.
- In order to shield hosts/end stations/ etc. from interpreting various proprietary BPDUs messages on the ring, an RPR specific TCN control message is recommended.
- The SRCs sublayer generates RPR TCN control message upon indication of its associated bridge management entity generating a STP TCN BPDU message.
- All station's receiving the RPR TCN control message advance their mapping table aging timers.
- The group still needs to determine and specify the method for the SRCs generating the RPR TCN control message.
  - Either an explicit control message from the bridge control entity to the SRCs.
  - SRCs snoops BPDUs transmitted by its bridge control entity.





## Station/Ringlet Selection in 802.17 MAC Model



Need to specify where the station ID mapping table resides in the 802.17 MAC model.

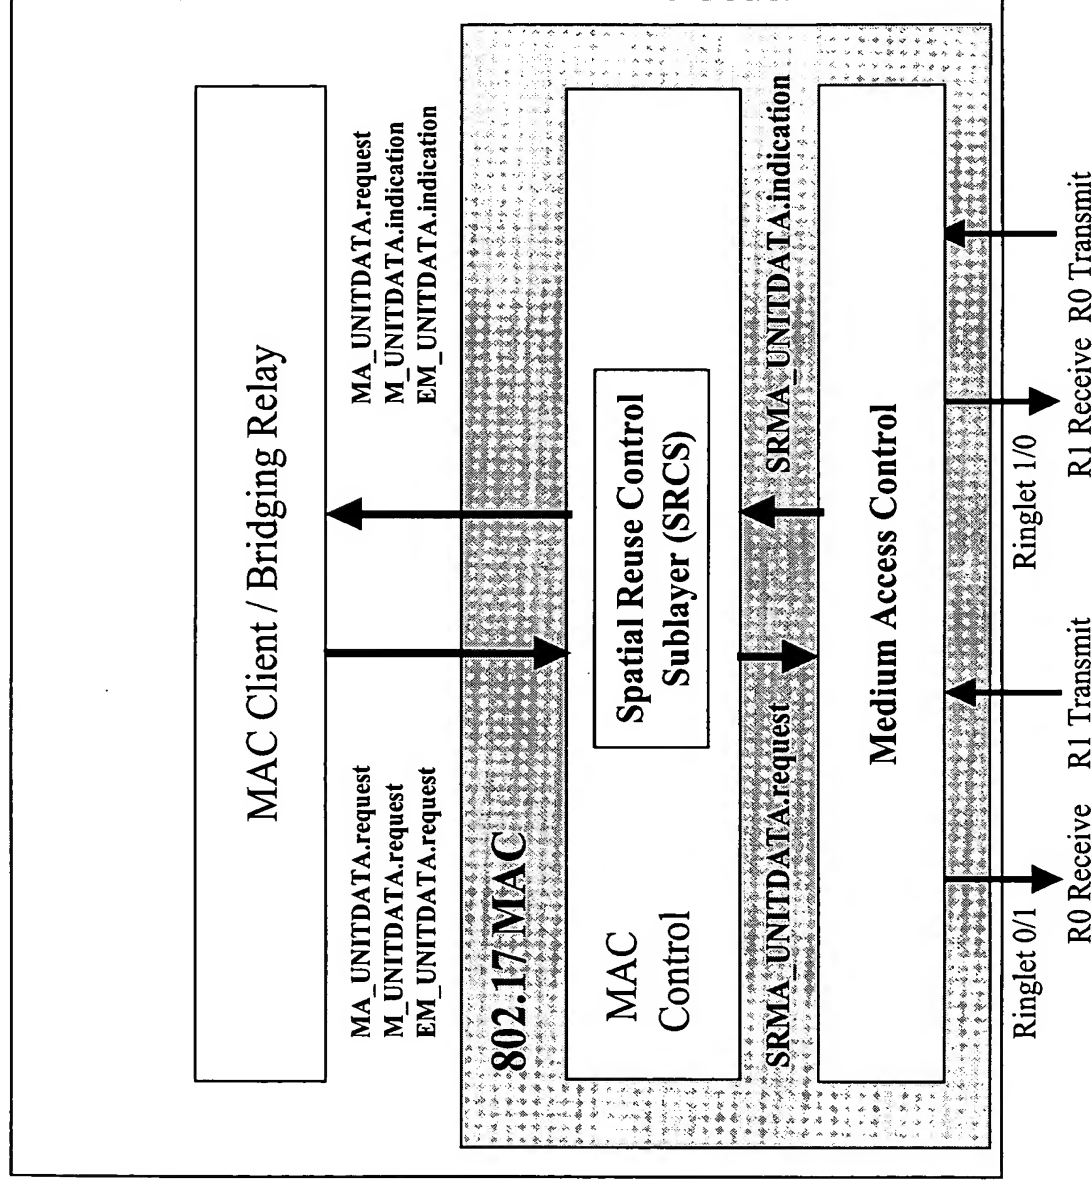
- Within the Medium Access Control
  - Within the MAC Control Sublayer (below MSAP)
  - Above the MAC Control Sublayer (above the MSAP)
- 
- Determining the sublayer should take into account
    - The rules for ringlet and station ID selection must be well specified in the 802.17 standard and adhere to the 802.1D/Q compatibility requirements.
    - The current draft optionally allows the MAC client to perform ringlet selection. The client shall also have to perform station ID selection for spatial reuse compatibility.

## Spatial Reuse Control Sublayer

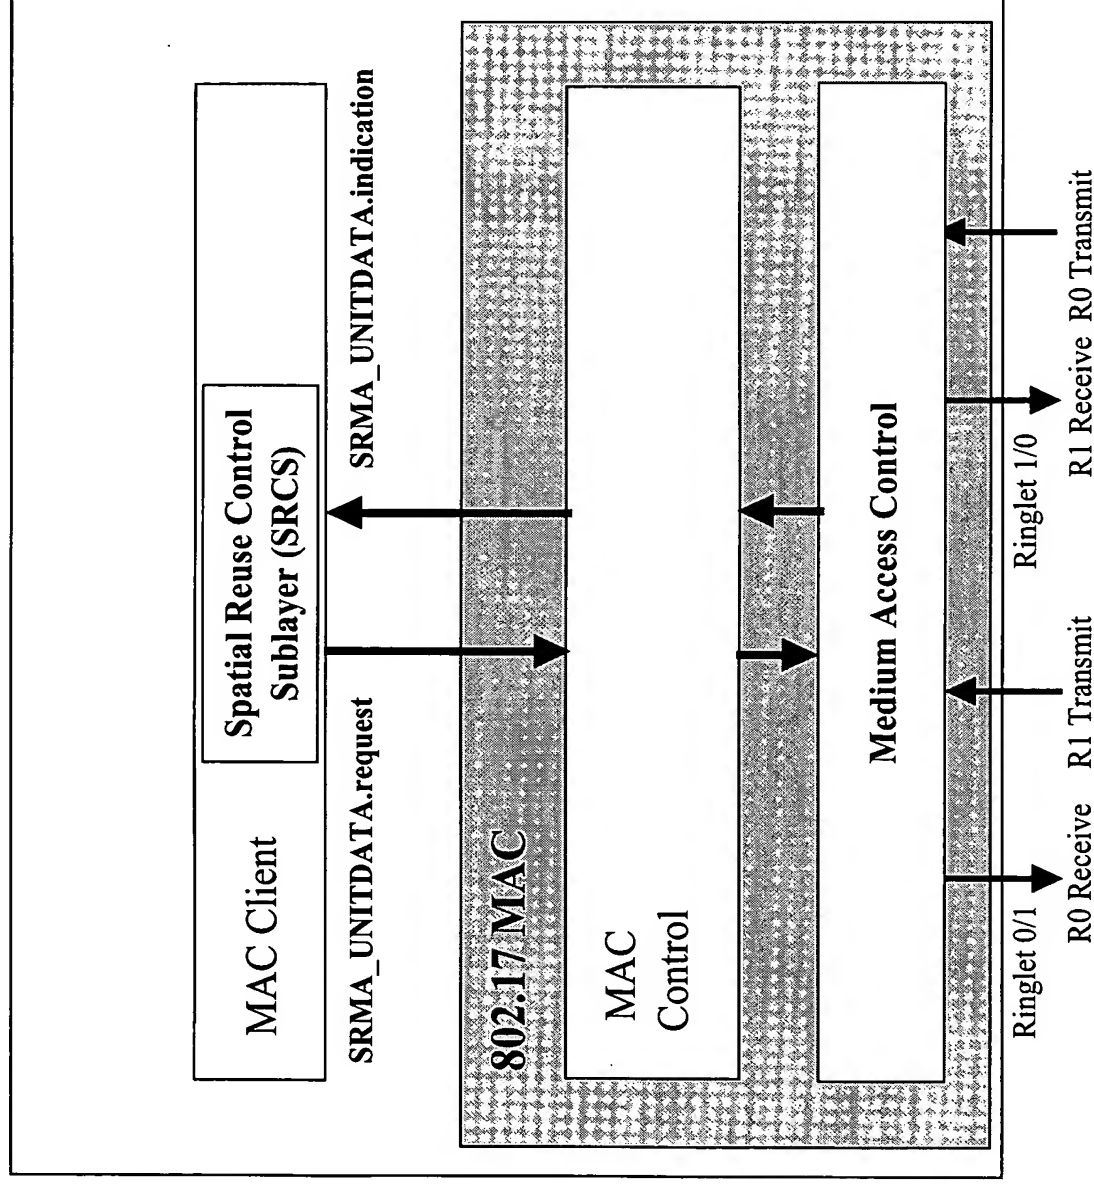



- Sublayer that implements the station and ringlet selection and adheres to the ringlet selection rules.
  - Supported by the SRMA\_DATA.request / indication primitives
  - SRMA\_DATA.request/indication primitives includes - Source/Destination MAC addresses, VLAN ID, ringlet ID, and source/destination station Ids.
  - Spatial Reuse Control Sublayer may either reside in the MAC or above the MAC depending on the entity responsible for making ringlet selection decisions in the station.
  - When residing above the MAC, the SRCs must conform to the same behavior as when operating within the MAC.

# 802.17 Unaware Client SRCS below the MSAP



# 802.17 Aware Client SRCS above the MSAP





The diagram illustrates the internal structure of a network device, showing the interaction between the SR Client and the MAC Control block.

**802.1 Bridge Relay Entity** (Forwarding, Filtering, Learning) is at the top, with *ISS* on the left and *JSS* on the right.

The **SR Client** block contains the **Spatial Reuse Control Sublayer (SRCS)**. It has two interfaces: **M\_UNITDATA.request** and **EM\_UNITDATA.request** on the left, and **M\_UNITDATA.indication** and **EM\_UNITDATA.indication** on the right.


The **802.17 MAC** block contains the **MAC Control** and **Medium Access Control** components. It has two interfaces: **SRMA\_UNITDATA.request** and **SRMA\_UNITDATA.indication** on the left, and **R0 Receive**, **R1 Transmit**, **Ringlet 0/1**, **Ringlet 1/0**, **R0 Receive**, and **R1 Transmit** on the right.



## Spatial Reuse Example



Unicast Frame Transmitted from C1 to C2, with frame marked as *flood = false*  
Ringlet selection in station 1 directs frame to intended destination  
Other Intermediate Stations do not these unicast frames to their bridge relay  
Frame Stripped at Destination Station based on Station ID or End\_Station MAC Address  
Since the only stations that see the frame are 1 and 4, reply may be along asymmetrical path






## Support for basic bridging compliance model



Station which does not implement the SRCs, can implement a simple bridging model for 802.1D/Q compliance. In this case, every bridge on the ring must see frames sourced by stations not implementing SRCs. The non-SRCS station must indicate `flood_always = true` for all transmitted frames. This signal indicates the frame is not to be destination stripped and all bridges attached to the ring must forward the frame to their bridging relay. Stations transmitting back to a station that has set its flood indicator must flood all frames intended for that station.

- Host A sets the `flood_always = true` indicator
- Frame is copied by all stations on the ring and frame is not destination stripped.
- All reply frames to host A are flooded.
- The `flood_always` indicator is needed in order for the SRCs receiving station to distinguish between transmitted packets from a SRCs station vs. a non-SRCS station. Unicast packets to a non-SRCS station are always flooded, whereas unicast packets to an SRCs station are always `DA_stripped`.





## Spatial Reuse for Multicast



- BAH consensus was to support spatial reuse for multicast traffic. The group discussed several different multicast SR proposals.
  - Multicast frames are destination stripped by the last station on the ring in a multicast group. The DSID or TTL in the frame may be used to indicate the last station in the group which is to destination strip the frame.
  - The MAC multicast address in the frame may also be used to perform additional multicast filtering. Some client's may prefer the MAC to filter received multicast traffic.
- Further work is required to complete the definitions for spatial reuse of multicast traffic.





## Further Work Items

- RPR Topology Control Messages  
Spatial Reuse for Multicast frames



## Conclusions



Spatial Reuse Control Sublayer defines a general method for bridges and end stations to both achieve spatial reuse and maintain 802.1D/Q compliance.

- Current draft frame format needs to be enhanced to support station Ids.
- Further work is needed to define operation of RPR TCN control messages
- Current draft has an 802.1D/Q compliance issue when stripping / flooding stations coexist on the ring. This issue must be addressed to achieve any form of spatial reuse in 802.17 and 802.1D/Q bridging compliance.
- Current draft does not support bridging nodes on the ring. Have we satisfied the PAR or internal WG requirements if the standard precludes spatial reuse of bridged traffic on the ring?
- Adopting the general method for spatial reuse obviates the need basic bridging compliance. A special flood bit in the header is needed for spatial reuse nodes to coexist with basic bridging compliance nodes.



# **Backup Slides**


## **Addressing Issues Raised by Norm**

## Issues by Norm




- The issues raised by Norm, with proper RPR wording, is presented here
- Solution to those issues requires
  - Flooding for address learning purposes
  - A ringlet selection algorithm so the selected ringlet will not change frequently to avoid frame misordering.
- How to “flood” and achieving “spatial reuse” at the same time is also presented. Currently, two mechanisms are available to achieve both objectives. In this case, the issues raised by Norm can be resolved, and spatial reuse objective of RPR working group is also achieved.


A good ringlet selection algorithm need to be worked out by the group




- Assumption
  - Y transmits to X via B2 on the CCW ringlet and via B3 on the CW ringlet
  - If X does not implement IVL then X learns Y on the CW ringlet (blue from B3) and then relearns Y on the CCW ringlet (red from B2).
- The Issue
  - Transmission from X to Y oscillates between B2 and B3 (whichever sent the last frame to X), and hence the data delivered will be out of sequence
- Solutions
  1. X performs flooding
  2. X performs static ringlet selection for a given MAC address
  3. X performs IVL ringlet selection when implementing Spatial Reuse. Others devices support same (MAC, VLAN, FID) mapping.



- Assumption
  - Y transmits to X via B2 on the CCW ringlet and via B3 on the CW ringlet
  - If X does not implement IVL then X learns Y on the CW ringlet (blue from B3) and then relearns Y on the CCW ringlet (red from B2).
  - X performs “steering only” protection option for all the frames for destination Y
- The Issue
  - Transmission from X to Y oscillates between CW and CCW (depending on whether last frame received by X was from B2 or B3), and hence the data delivered will be lost. CW transmission loses blue traffic, and CCW transmission loses red traffic.
- Solutions
  1. X performs flooding.
  2. X performs static ringlet selection with wrapping protection
  3. X performs IVL ringlet selection when implementing Spatial Reuse.




- **Assumption**
  - Unidirectional VLANs between X and Y. X transmits to Y on blue, Y transmits to X on red VLAN
  - X implements IVL
- **The Issue**
  - X never learns Y MAC address on the blue VLAN
  - X always floods Y.blue traffic on the ring.
  - B2 never sees Y MAC address and hence always floods the frame to all its interfaces (as shown), regardless how it is configured
- **Solutions**
  1. X and B3 perform flooding. B2 bridge relay needs to support SVL to prevent excessive flooding.
  2. X performs SVL ringlet selection. All other bridge relays support same SVL (MAC, VLAN, FIB) mapping to prevent excessive flooding




- Assumption
  - Old path is physically broken and a new path, after spanning tree convergence, is selected
- The Issue
  - TCN propagated upstream along new path. If X is a router will not participate in TCN
  - X continues to transmit frames to Y through B3 if Y is not transmitting. Packets are blackholed.
- Solutions
  1. X performs flooding.
  2. X processes TCN STP messages as part of ringlet selection for spatial reuse.





- Assumption
  - Asymmetrical traffic flows
  - In 802.17 protection scenario where (i.e. a node between B3 and X is bad, X is using “wrapping” and B3 is using “Steering”)
- The Issue
  - B2 never sees address X, resulting un-necessary flooding on other ports of B2.
- Solution
  - X performs flooding
  - X implements an address updating mechanism, such that other stations on the ring can learn X.

# Sixth Issue (Spatial Reuse mode and Promiscuous mode co-exist on the same ring)



- Assumption
  - For some reason, the traffic is uni-directional (I.e. a node between B3 and X is bad, X is using “wrapping” and B3 is using “Steering”)
  - B3 is flooding all the time
- The Issue
  - B2 will never see address X, resulting un-necessary flooding on other ports of B2, for all the Y to X frames.
  - This is not much different from the fifth issue
- Solution
  - X performs flooding
  - X implements an address updating mechanism; such that other stations on the ring can learn X.



# Thank You!!